Interrupts, edge- and level-triggered

NOTE: This article is based on old information and needs to be rewritten.

Edge-triggered interrupts are identified by a simple and transient voltage changed on the interrupt lines as adapters request service from the device drivers.

This method of signaling has several limitations. The programmable interrupt controller not only can miss this instantaneous voltage transition, but also often misinterpret noise on the lines as valid interrupts. The system and adapter circuitry, as well as the fundamental signaling nature of edge-triggered interrupts, made interrupt sharing impossible on ISA systems, where most adapters only support edge-triggered interrupts.

To address these limitations, both EISA- and MCA-based systems equipped with sophisticated interrupt controllers provide level-sensitive interrupts. Unlike signal or voltage transitions in edge-triggered systems, level-triggered interrupts are marked by not only a transition, but also continuation of the new voltage until the adapter drops the signal on instructions from the device driver.

While all MCA adapters use level-sensitive interrupts, and the operating system programs the PIC to reflect this, EISA interrupts can be individually programmed as either edge or level, and the driver package should reflect its preference by marking its `itype` value as 4 in the *sdevice* system file. The use of `itype` in this instance is overloaded in that the driver not only indicates a willingness to share interrupts with other device/driver combinations, the kernel automatically programs the PIC to mark the interrupt as level-triggered.

NOTE: If the hardware is programmed for level-triggered interrupts (for example, using an EISA configuration utility), the driver should inform the kernel by specifying the appropriate value in the *sdevice* file.

The kernel cannot and will not cross-verify the device and PIC level interrupt settings.

HDK 7.1.0 - June 1999